

FIELD GUIDE

Choose the work.

A field guide to a useful first AI initiative.



A NOTE TO THE READER

A useful starting point.

The first decision is what deserves attention. A new capability can create interest without resolving a real operating problem. Start with the work people are already trying to complete, then ask where a better system could change the outcome.

This guide is designed for a working conversation between a business sponsor, a process owner, and the people who understand the relevant information and systems. Bring a recently completed task and the material used to carry it out. Leave with a bounded opportunity and a decision you can test.

The questions are prompts for your own judgment. They do not prescribe an investment, a delivery schedule, or a guaranteed return.

IN THIS EDITION

Follow one piece of work.	3
Choose the right kind of help.	4
Define a boundary that matters.	5
Make the next decision testable.	6

CHOOSE THE WORK.

Follow one piece of work.

Look for the task behind the request for technology.

Choose a task that has actually reached completion. Follow it from the first signal to the accepted result. Include the messages, records, checks, and handoffs that made the result possible. The process described in a meeting can differ from the process people use when a request arrives with missing information.

Separate hands-on effort from waiting. A faster draft may be useful, but it will not remove a queue for approval. Record where the task returns for clarification and which person decides that it is complete. Notice whether another team absorbs correction work after the visible handoff.

Use several ordinary cases and at least one difficult case before drawing a conclusion. The objective is a shared picture of the work, including its exceptions. A clear map can reveal an improvement that requires a simpler rule, a better source, or a clearer owner before it requires AI.

QUESTIONS TO TAKE INTO THE ROOM

What event begins the task, and what observable result ends it?

Where do people search, repeat, wait, or ask for clarification?

Who is responsible when the task does not follow the ordinary path?

CHOOSE THE WORK.

Choose the right kind of help.

Distinguish movement, rules, and interpretation.

Moving an approved value between systems is different from deciding what an ambiguous message means. Mark the steps that transfer information, apply a stable rule, and interpret varied material. Conventional integrations can handle much of the first category; explicit business rules often handle the second.

AI becomes a candidate where interpretation is useful and someone can assess whether the result is good enough. A draft, a classification, or a source-grounded summary can be a useful contribution inside a larger workflow. It does not need authority to make every downstream decision.

If experienced people disagree on the correct outcome, understand why. They may be using different evidence, exercising legitimate discretion, or working around an unsettled policy. Document the distinction. An implementation should not quietly choose a business rule that the organization has not agreed.

QUESTIONS TO TAKE INTO THE ROOM

Which steps can be expressed as a stable rule?

Where would interpretation add value, and how would a reviewer check it?

Which decisions must remain with an accountable person?

CHOOSE THE WORK.

Define a boundary that matters.

Keep the first scope small enough to observe and useful enough to evaluate.

Write one input, one intended output, and the point of review. For a request-handling workflow, the initial system might assemble relevant context and prepare a response for an experienced reviewer. That is a different commitment from sending a response, changing a record, or promising a commercial outcome.

Name the sources the system may use and the actions it may take. Decide what it should do when an input is missing, documents conflict, or an integration is unavailable. An explicit pause with an owner and a next action is more useful than a plausible answer that hides missing evidence.

Keep a workable manual route. The first version should be able to stop without making unfinished tasks disappear. Before adding another integration or expanding the audience, ask whether the existing boundary has answered the question it was intended to test.

QUESTIONS TO TAKE INTO THE ROOM

What is the smallest complete output that helps the team?

What can the system read, prepare, change, or send?

How does a person recover the task if the system stops?

CHOOSE THE WORK.

Make the next decision testable.

Agree on evidence before becoming attached to an implementation.

Write the decision the initiative should resolve: continue, narrow the scope, change the approach, or stop. Assemble representative cases and describe a useful result for each. Include missing context, ambiguous inputs, and cases that should be escalated. Keep some cases separate from those used to improve the system.

Compare the whole task under the same definition of completion. Count preparation, checking, corrections, and follow-up. Record important mistakes separately from stylistic edits. An average result should not conceal a failure that makes the workflow unsuitable for its intended authority.

At the review point, examine the evidence with the operational owner. A useful outcome can be a narrower workflow or a decision that the information is not ready. If the next step is live use, define ownership, access, monitoring, and recovery before increasing exposure.

QUESTIONS TO TAKE INTO THE ROOM

What evidence would change our decision?

How will we count review and correction effort consistently?

What would cause us to narrow the scope or stop?

KEEP THE CONVERSATION USEFUL

The next considered step.

Bring one real task, its owner, and the source material into the room. A useful starting point is a clear workflow, a deliberate boundary, and a question the work can answer.

FURTHER READING

[Anthropic: Building effective agents](#)

Further reading on choosing between predefined workflows and model-directed agents. Tooling references change; use current provider documentation for implementation.

<https://www.anthropic.com/engineering/building-effective-agents>

[OpenAI: Evaluation best practices](#)

Further reading on task-specific evaluation, representative cases, and continuous review.

<https://developers.openai.com/api/docs/guides/evaluation-best-practices>

Written by Ayan. First edition, October 2026. This publication is general planning guidance. Project scope, fees, timing, and responsibilities are agreed separately in writing.

withayan.com / hello@withayan.com